How Fast do Identifier-to-locator Mappings Change in Networks with Identifier/Locator Separation?

Hongbin Luo, Hongke Zhang
School of Electronic and Information Engineering
Beijing Jiaotong University, Beijing 100044, China
Email: {hbluo,hkzhang}@bjtu.edu.cn

Chunming Qiao
Department of Computer Science and Engineering
SUNY at Buffalo, Buffalo, NY 14260-2000, U.S.A.
Email: qiao@computer.org

***Abstract*—There is a growing consensus that identifier/locator separation is a promising solution to the scalability issue of the current routing infrastructure. After locators are separated from identifiers, end hosts roam from place to place without changing their identifiers. On the other hand, their locators change when they roam from one place to another, which leads to changes of identifier-to-locator mappings. In this paper, we analyze the possible change rates of identifier-to-locator mappings by analyzing the intervals of consecutive handovers, based on real data traces of collected from 2,348 buses, 536 taxis, and several tens of pedestrians. We believe that our results lay a solid foundation for the design and performance evaluation of mapping mechanisms that map identifiers onto locators in networks with identifier/locator separation.**

***Index Terms*—Routing, identifier/locator separation, identifier-to-locator mapping, mapping update rate.**

I. INTRODUCTION

In recent year, it is widely recognized that the current routing infrastructure faces a serious scaling issue [1]. In addition, there is a growing consensus that identifier/locator separation is a promising solution for this scaling issue [2] - [7]. In identifier/locator separation, the identity role of the current Internet Protocol (IP) addresses is represented by an identifier (ID) namespace, and the location role of the current IP addresses is represented by a locator namespace. In addition, identifiers are used in the application and transport layers for identifying nodes, and locators are used in the network layer for locating nodes in the network topology. This makes it possible for nodes to change locators at any time without disrupting ongoing communication sessions, thus supporting efficient mobility, and multi-homing.

After identifiers are separated from locators, in order to send packets to a host B, a host A generally needs to know the locator of host B so that the routing infrastructure knows where to send these packets. Since hosts roam from place to place, however, their locators change accordingly. Thus their identifier-to-locator mappings also change from time to time. As a result, a mapping system is required to map identifiers onto their current locators. For example, [8] - [11] proposed a set of mechanisms to map identifiers onto locators for Locator/Identifier Separation Protocol (LISP) [2]. The work in [12] further proposed an approach to map flat identifiers onto locators.

While various approaches including above mentioned have been proposed and have their pros and cons, a common drawback of them is that they mainly focus on the design of the mechanisms and lack detailed performance analysis. Based on real data traces, the work in [13] analyzed the possible map-request rates, assuming that identifiers are IP-alike aggregable. The work in [12] further analyzed the possible map-request rates, assuming that identifiers are flat. However, these results are based on data collected in the current Internet, in which end hosts rarely roam. Even worse, these work did not consider the overhead caused by the change of identifier-to-locator mappings due to the lack of knowledge on the rates that identifier-to-locator mappings will change, although there are some related work that focus on the mobility pattern of human beings [14], [15].

In this paper, we analyze the rates that identifier-to-locator mappings may change by analyzing the intervals of consecutive identifier-to-locator changes. For this purpose, we collect real traces of 2,348 buses at Shanghai, 536 Taxis at San Francisco, and several tens of pedestrians at five other places. All buses, taxis and pedestrians are equipped with GPS receivers that record their positions in terms of (latitude, longitude) and the time they appeared in their positions. Notice that we do not use real data traces of mobile phone users because only when a mobile user initiates or receives a call or a text message, the location of the tower routing the communication (instead of the location of the mobile user) is recorded. Since mobile users cannot communicate with others all the time, the data traces for them is significantly less accurate than these collected from GPS receivers that record their current position every one or ten seconds.

In order to analyze the intervals of consecutive identifier-to-locator changes, we use a set of squares with area of $S\ km^2$ to cover the traveling area of each bus, taxi, or pedestrian. As a result, buses, taxis, or pedestrians travel from one square into another, stay some time and then move to a new one. We thus record the time that a bus, taxi, or pedestrian moves into a new square and let the interval of consecutive times be the interval of consecutive handovers. Notice that we use squares (instead of circles or rectangles) to cover the traveling area of a bus, taxi, or pedestrian for computation simplicity. In addition, since we care about subnet instead of base-stations, it is practical for subnets to covers square areas.

From the data analysis of our traces, we find the followings:

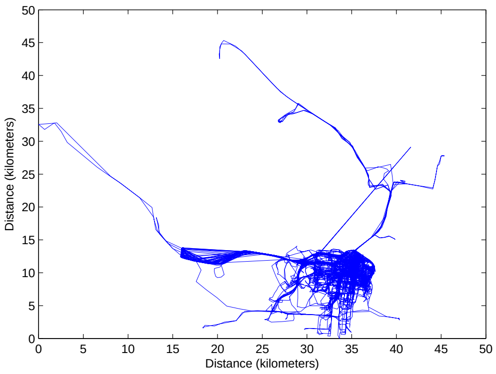

Fig. 1. Trace of a randomly selected taxi.

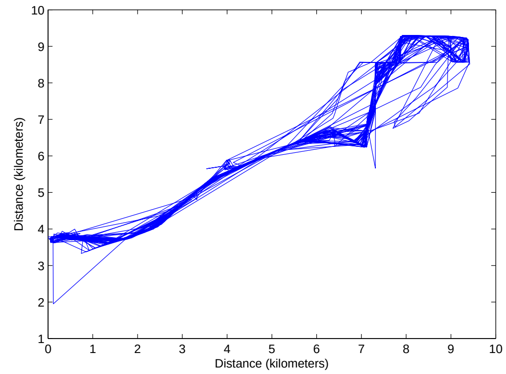

Fig. 2. Trace of a randomly selected bus.

1) For pedestrians, with a probability higher than 99%, the intervals between two consecutive handovers are less than 21,690 and 34,200 seconds when S are 1 km^2 and 3 km^2, respectively. In addition, with a probability higher than 90%, the intervals between two consecutive handovers are less than 3,000 and 4,800 seconds when S are 1 km^2 and 3 km^2, respectively. 2) For taxis, with a probability higher than 99%, the intervals between two consecutive handovers are less than 2,120 and 3,360 seconds when S are 1 km^2 and 3 km^2, respectively. In addition, with a probability higher than 90%, the intervals are less than 310 and 527 seconds when S are 1 km^2 and 3 km^2, respectively. 3) For buses, with a probability higher than 99%, the intervals between two consecutive handovers are less than 3,200 and 6,500 seconds when S are 1 km^2 and 3 km^2, respectively. In addition, with a probability higher than 90%, the intervals are less than 530 and 980 seconds when S are 1 km^2 and 3 km^2, respectively.

The rest of the paper is organized as follows. In Section II, we describe how data traces are collected and how to analyze the collected data. In Section III, we present the results from our analysis. Finally, we conclude the paper in Section IV.

II. Data Collection and Analytical Method

In this section, we first present how data are collected. We then describe the methodology used to analyze the interval between consecutive handovers.

A. Data Collection

For our purpose, we collect three types of data. The first type is the traces of 536 taxis at San Francisco. The second type is the traces of 2,348 buses at Shanghai, and the third type is several tens of pedestrians at five different sites. We use these three types of data traces because we often go to a further place either by bus or by taxi, and to a nearby place on foot in our daily life. In addition, taxis run faster than buses since buses frequently stop at some bus stops but taxis do not need to stop at such bus stops. Furthermore, it is evident that buses run faster than we walk. As a result, choosing these types of data traces makes it possible for us to analyze the intervals between two consecutive handovers for different mobility patterns.

In order to collect the traces of buses and taxis, we place GPS receivers at them that are capable with a position accuracy of better than three meters 95% of the time. The GPS receiver at each taxi or bus takes reading of its current positions in terms of (latitude, longitude) at every second, and records them into a track log. For each taxi, we record its trace of 24 continuous days, which results in a total traces of 12,864 days. Figure 1 shows a trace randomly selected from the 536 traces. From this figure, we observe that the taxi travels in an area with about 50 kilometers length and about 50 kilometers width. However, at the most majority of the time, the taxi travels in an area with about 25 kilometers length (from about 15 to about 40 kilometers in Figure 1) and 15 kilometers width (from about 0 to about 15 kilometers in Figure 1).

For each bus, we record its trace for five continuous days, which results in a total traces of 11,740 days. Figure 2 shows a trace randomly selected from the 2348 traces. From this figure, we observe that the selected bus travels along a given path in most cases, and occasionally travels to some other places (*e.g.*, go to the gas station). In addition, since a bus often travels along a given path, its traveling area is significantly smaller than that of a taxi. For example, the taxi producing the trace shown in Figure 1 travels in an area of about 2,500 km^2 but the bus producing the trace shown in Figure 2 travels in an area of about only 100 km^2.

The traces for pedestrians are borrowed from [14]. In particular, five sites are chosen. These include two university campuses (NCSU and KAIST), New York City, Disney World at Orlando, and North Carolina State Fair. In order to collect data, Garmin GPS 60CSx handheld receivers that are wide area augmentation system capable with a position accuracy of better than three meters 95% of the time are used. In addition, every GPS receiver takes reading of its current position at every 10 seconds and record them into a daily track log. 44 different pedestrians participate the data collection and produce 226 daily traces. Figure 3 shows a trace randomly selected from the 226 traces. From Figure 3, we observe that the pedestrian produced the trace shown in Figure 3 walks around an area with a length of only 1.2 kilometers and a width of about 1.2 kilometers.

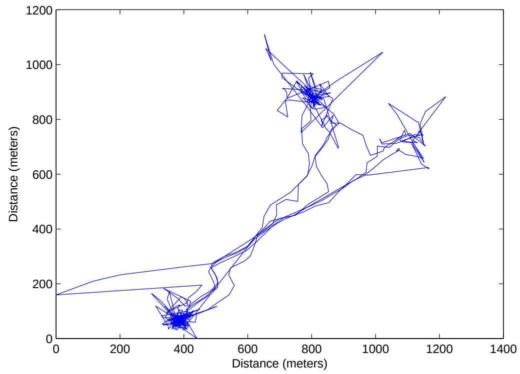

Fig. 3. Trace of a randomly selected pedestrian.

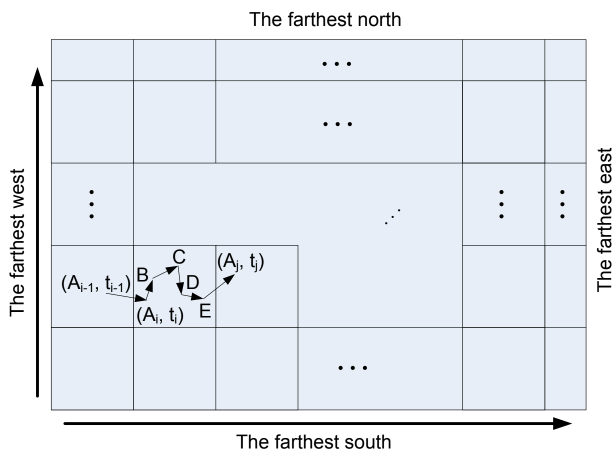

Fig. 4. Layout of subnets used for analyzing intervals between consecutive handovers.

B. Analytical Method

We analyze our data in two steps. In the first step, we map the traces into a two-dimensional area since the GPS receivers produce three-dimensional positions. For the traces of pedestrians, we further recompute a position at every 30 seconds by averaging the three samples over that 30 second period [14].

In the second step, we analyze the intervals between consecutive handovers. For this purpose, we first place each trace into a rectangular topology defined by the farthest east, the farthest north, the farthest west, and the farthest south of the trace, as illustrated in Figure 4. We then divide the rectangular topology into a set of squares, each of which has an area of S km^2, thus a width of $\sqrt{S}$ kilometers. Notice that a rectangular topology may not be divided into multiple squares with area of S km^2. To deal with this issue, we place some squares in the rectangular topology from the farthest south and the farthest west to the farthest north and the farthest east, until the rectangular topology cannot place squares with area of S km^2. In addition, we divide the rest area into some small rectangles with width of $\sqrt{S}$ kilometers and/or length of the rest, as illustrated in Figure 4. As a result, every record in a

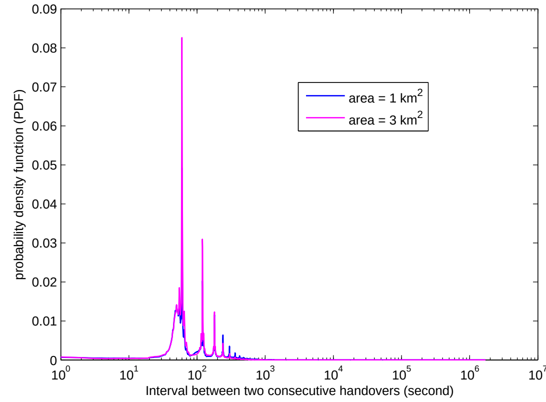

Fig. 5. Probability density function of handover intervals between handovers for taxis.

trace must correspond to a point in the rectangular topology corresponding to the trace. This way, every square or a small rectangle corresponds to a subnet, and when a mobile user moves from one subnet into another one, a handover happens.

We count the intervals between consecutive handovers as follows. We first place the first record in a trace into the rectangular topology corresponding to the trace and record the time corresponding to the first record as the time a mobile user moves into a subnet. After that, every time the mobile user moves into a new subnet, we record the time the user enters into the new subnet until the whole trace ends. We then compute the interval of two consecutive records as the interval of consecutive handovers. For example, when the user shown in Figure 4 moves to point A_i at time t_i, we record t_i since, at time t_i, the user moves into a new subnet from the subnet at which point A_{i-1} locates. After some time when the user passes through points B, C, D, and E, the mobile user moves to point A_j at time t_j when the user moves to a new subnet. As a result, we also record t_j and the interval between time t_j and time t_i as the interval of two consecutive handovers.

III. Results

Using the method described in the above section, we are able to get a set of intervals between consecutive handovers. Based on these intervals, we are able to plot their probability density functions and cumulative density functions. In particular, we consider the two cases that S are 1 km^2 and 3 km^2 because of two reasons. First, almost all power towers in mobile networks covers 1 or 3 km^2. Second, the population density in most cities ranges from several thousands to several tens of thousands. For example, the population densities in Beijing and Shanghai are 11,500 and 13,400 per km^2, respectively. At the same time, most routers in current edge networks support at most 50,000 users, which implies that a subnet with a single edge router is only able to cover an area of 3 km^2.

In Figure 5, we plot the probability density function of handover intervals for taxis when S is 1 km^2 and 3 km^2.

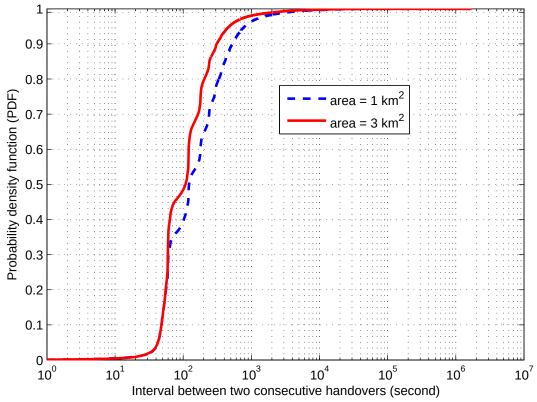

Fig. 6. Cumulative density function of handover intervals between handovers for taxis.

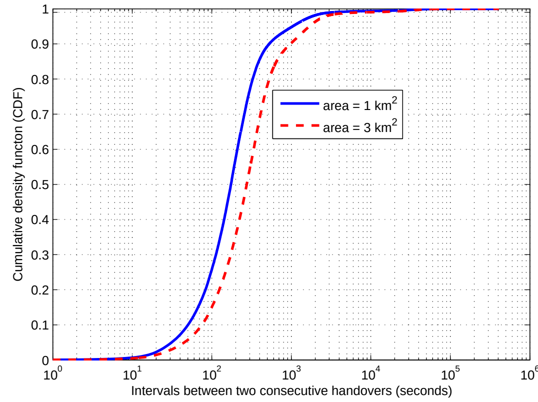

Fig. 8. Cumulative density function of handover intervals between handovers for buses.

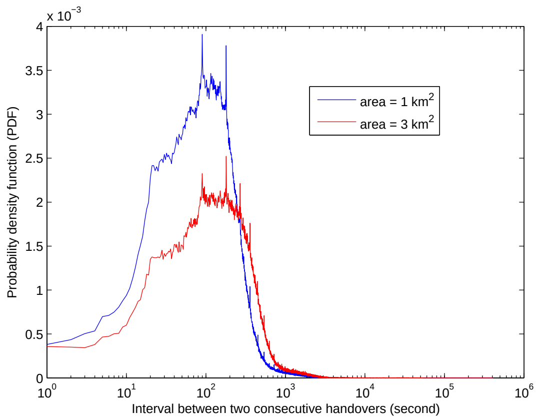

Fig. 7. Probability density function of handover intervals between handovers for buses.

From this figure, it is clear that most intervals are between 20 seconds and 300 seconds for both cases. In particular, we observe that the handover interval of 60 seconds has the largest probability of about 0.082. This may correspond to the case the a taxi passes through a subnet in a speed of about 60 kilometers per hour which is a normal speed of taxis at San Francisco. In addition, we also observe that the intervals between consecutive handovers when S equals to 1 km^2 are smaller than those when S equals to 3 km^2. This is because when the size of a single subnet is larger, it spends more time for a taxi to pass through the subnet. Thus the interval between two consecutive handovers becomes larger. Notice that there are some very long intervals (*e.g.*, 10^6 seconds). This is because some taxis do not put into service in such intervals.

Figure 6 further shows the cumulative density function of intervals between consecutive handovers for taxis when S are 1 and 3 km^2. From this figure, we further observe that the interval between consecutive handovers when S is 1 km^2 is smaller than that when S is 3 km^2. More importantly, we observe that, with a probability higher than 99%, the intervals between two consecutive handovers are less than 2,120 and 3,360 seconds when S are 1 km^2 and 3 km^2, respectively. In addition, with a probability higher than 90%, the intervals between consecutive handovers are less than 310 and 527 seconds when S are 1 km^2 and 3 km^2, respectively.

In Figure 7, we show the probability density function of handover intervals between consecutive handovers for buses when S is 1 km^2 and 3 km^2. From this figure, it is clear that the interval between consecutive handovers when S is 3 km^2 is longer than the interval when S is 1 km^2. In addition, we also observe that most intervals are smaller than 2,000 seconds for both S = 1 km^2 and S = 3 km^2. When compared with the probability density function for taxis, we observe that the probability of having smaller intervals for buses is higher than that for taxis. Indeed, when S = 1 km^2, 1% intervals are less than 22 seconds and 13 seconds for taxis and buses, respectively. In addition, when S = 3 km^2, 1% intervals are less than 21 seconds and 17 seconds for taxis and buses, respectively. Furthermore, 30% intervals for taxis are less than 60 seconds. By contrast, less than 20% intervals for buses are less than 60 seconds.

We further plot the cumulative density function of intervals between consecutive handovers for buses in Figure 8. As shown, with a probability higher than 99%, the intervals between consecutive handovers are smaller than 3,200 seconds and 6,500 seconds when S are 1 km^2 and 3 km^2, respectively. In addition, with a probability higher than 90%, the intervals between consecutive handovers are smaller than 530 seconds and 980 seconds when S are 1 km^2 and 3 km^2, respectively. Comparing the cumulative density functions of intervals for taxis and buses, we also observe that the intervals for buses are smaller than that for taxis. This may be because buses often stop at bus stops along their paths and often travel along fixed routes. Another reason for this observation is that taxis often run faster than buses.

Figure 9 plots the probability density function of intervals between consecutive handovers for pedestrians when S are 1 and 3 km^2. From this figure, we observe that the probability

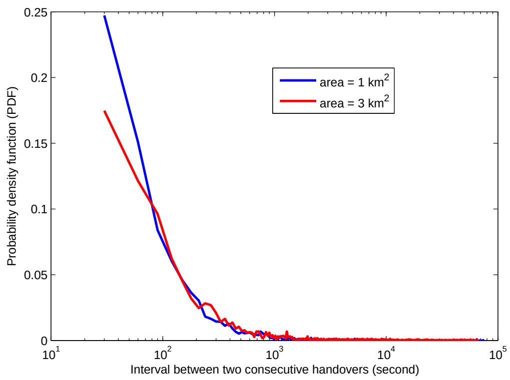

Fig. 9. Probability density function of handover intervals between handovers for pedestrians.

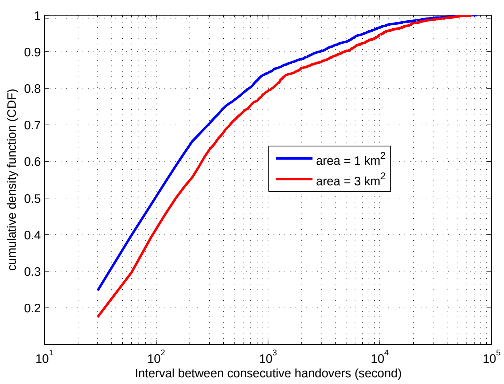

Fig. 10. Cumulative density function of handover intervals between handovers for pedestrians.

generally decreases with the increase of intervals for both S = 1 km^2 and S = 3 km^2. In particular, the probabilities that the intervals between consecutive handovers are 30 seconds for both S = 1 km^2 and S = 3 km^2 are of about 25% and 18%, respectively. The reason for this difference is that when the area of a single subnet is smaller, it is more likely that a pedestrian walks across different subnets. For example, for the pedestrian produced the trace shown in Figure 3, the pedestrian traveled only within an area of about 1.68 km^2. As a result, if a subnet covers an area of 1 km^2, the pedestrian should produce some handovers. On the other hand, if a subnet covers an area of 3 km^2, the pedestrian cannot produce any handover.

In Figure 10, we further plot the cumulative density function of intervals between consecutive handovers for pedestrians when S are 1 and 3 km^2. From this figure, we observe that, with a probability higher than 99%, the intervals between consecutive handovers are less than 21,690 seconds and 34,200 seconds when S are 1 km^2 and 3 km^2, respectively. In addition, with a probability higher than 90%, the intervals between consecutive handovers are less than 3,000 seconds and 4,800 seconds when S are 1 km^2 and 3 km^2, respectively. Comparing the results for pedestrians, buses, and taxis, it is apparent that the intervals for pedestrians are longer than those for buses and taxis.

IV. Conclusion

In this paper, we have investigated the possible change rates of identifier-to-locator mappings, based on real data traces collected from taxis, buses, and pedestrians. We have found that, depending on different mobility patterns (by bus, by taxi, or walk), the distributions of the interval of consecutive handovers are different. We have also found that, with probabilities higher than 99% and 90%, the intervals between consecutive handovers are less than 34,200 seconds and 4,800 seconds, respectively, in real applications.

Acknowledgment

We would like to thank Prof. Seongik Hong from North Carolina State University for providing the data traces of pedestrians. This work is supported in part by the "973 Program" of China under Grant No. 2007CB307100, and in part by the Natural Science Foundation of China (NSFC) under Grant No. 60833002 and 60702008.

References

[1] D. Meyer, L. Zhang, and K. Fall, "Report from the IAB workshop on routing and addressing," *IETF Internet Standard, RFC 4984*, Sep. 2007.

[2] D. Farinacci, V. Fuller, D. Oran, D. Meyer, and S. Brim, "Locator/ID separation protocol (LISP)," *IETF Internet Draft*, draft-ietf-lisp-05.txt (work in progress), Sep. 2009.

[3] D. Massey, L. Wang, B. Zhang, and L. Zhang,"A Scalable Routing System Design for Future Internet," in *Proc. of ACM International Workshop on IPv6 and the Future of the Internet (IPv6)*, Kyoto, Japan, Aug. 2007.

[4] C. Vogt, "Six/One Router: A Scalable and Backwards Compatible Solution for Provider-Independent Addressing," in *Proc. of ACM SIGCOMM MobiArch'08*, Seattle, WA, USA, Aug. 2008.

[5] D. Jen, M. Meisel, D. Massey, L. Wang, B. Zhang, and L. Zhang, "APT: a practical transit mapping service," *IETF Internet Draft*, draft-jen-apt-01.txt, Nov. 2007.

[6] X. Xu and R. Jain, "Routing Architecture for the Next Generation Internet (RANGI)," *IETF Internet Draft*, draft-xu-rangi-03.txt (work in progress), Feb. 2010.

[7] R. Whittle, "Ivip (Internet vastly improved plumbing) architecture," *IETF Internet Draft*, draft-whittle-ivip-arch-04.txt (work in progress), March 2010.

[8] E. Lear, "NERD: A Not-so-novel EID to RLOC Database," *IETF Internet Draft*, draft-lear-lisp-nerd-08.txt (work in progress), March 2010.

[9] S. Brim, N. Chiappa, D. Farinacci, V. Fuller, D. Lewis, and D. Meyer, "LISP-CONS: A content distribution overlay network service for LISP," *IETF Internet Draft*, draft-meyer-lisp-cons-04.txt, April 2008.

[10] D. Farinacci, V. Fuller, and D. Meyer, "LISP alternative topology (LISP+ALT)," *IETF Internet Draft*, draft-ietf-lisp-alt-02.txt (work in progress), Jan. 2010.

[11] L. Mathy, L. Iannone, and O. Bonaventure, "LISP-DHT: towards a DHT to map identifiers onto locators," *IETF Internet Draft*, draft-mathy-lisp-dht-00.txt, Feb. 2008.

[12] H. Luo, Y. Qin, and H. Zhang, "A DHT-based identifier-to-locator mapping approach for a scalable Internet," *IEEE Transactions on Parallel and Distributed System*, vol. 20, no. 12, Dec. 2009.

[13] L. Lannone, and O. Bonaventure, "On the cost of caching locator/ID mappings," in *Proc. of ACM CoNEXT 2007*, NY, USA, Dec. 2007.

[14] I. Rhee, M. Shin, S. Hong, K. Lee, and S. Chong, "On the Levy-walk nature of human mobility," in *Proc. IEEE INFOCOM 2008*, Phoenix, AZ, USA, April 2008.

[15] M. C. Gonzalez, A. A. Hidalgo, and A.-L. Barabasi, "Understanding individual human mobility patterns," *Nature*, vol. 453, no. 5, June 2008, pp. 779 - 782.